

Amendments to the Claims:

This listing of claims will replace all prior versions, and listings of claims in the application:

Listing of Claims:

- 1 1. (Currently amended): A method for accessing a storage system
2 comprising:
3 accessing a data object, the data object being divisible into one or more partitions,
4 the partitions comprising data from the data object, the partitions referred to as input partitions;
5 and
6 for each input partition, if there are no other partitions among other data objects in
7 the storage system that are identical to the input partition, then producing one or more replicas of
8 the input partition.
- 1 2. (Original): The method of claim 1 wherein the data object is a file.
- 1 3. (Original): The method of claim 1 wherein a first partition is associated
2 with a partition ID, such that another partition having content that is identical to content of the
3 first partition, then the other partition is associated with the same partition ID as that of the first
4 partition, the method further comprising storing partition identification information comprising a
5 plurality of partition IDs, each partition ID being associated with one or more partitions,
6 determining a partition ID of the input partition, accessing the partition identity information to
7 determine if there are any partitions that are identical to the input partition based on the partition
8 ID of the input partition.
- 1 4. (Original): The method of claim 1 wherein the data object is to be stored
2 on the storage system, the method further comprising receiving a request to store the data object
3 on the storage system, receiving data comprising the data object, and storing the data object on
4 the storage system.

1 5. (Original): The method of claim 7 further comprising, for each input
2 partition, generating a content-based identifier based on at least some content of the input
3 partition and identifying first partitions in the storage system that have the same content-based
4 identifier, wherein the one or more replicas are produced if none of the first partitions is identical
5 to the input partition.

1 6. (Original): The method of claim 5 wherein the step of generating includes
2 applying a hash algorithm to at least a portion of the content of the input partition.

1 7. (Original): The method of claim 7 wherein the data object is a file.

1 8. (Currently amended): A method for accessing a storage system
2 comprising:
3 receiving data for a first file, to be stored in the storage system;
4 providing partition data from the first file which constitutes a first partition of the
5 first file;

6 if a number of second partitions in the storage system is less than a first
7 predetermined value, then producing a number of replicas of the first partition sufficient to
8 increase the number of second partitions to a second predetermined value, wherein each second
9 partition comprises data ~~that~~ belonging to a file in the storage system and is identical to the
10 partition data; and

11 if the number of second partitions is greater than a third predetermined value and
12 if there are one or more replicas of the first partition, then deleting one or more of the replicas,
13 wherein the number of second partitions is reduced; and

14 repeating for additional partition data comprising the first file.

1 9. (Original): The method of claim 8 further comprising receiving a request
2 to store the first file, receiving data comprising the first file, and storing the first file on the
3 storage system.

1 10. (Original): The method of claim 8 wherein partitions each is identified by
2 a content-based code and a group ID, wherein if data one partition is different from data of
3 another partition and both partitions have the same content-based code, then each partition has a
4 different group ID, whereby partitions that contain identical data are identified by the same
5 content-based code and group ID.

1 11. (Original): The method of claim 10 further comprising storing partition
2 identity information comprising a content-based code and a group ID that correspond to each
3 partition on the storage system, wherein the first partition is associated with a first content-based
4 code value and a first group ID value, wherein the number of second partitions can be
5 determined by consulting the partition identity information and counting the number of partitions
6 whose corresponding content-based code is equal to the first content-based code value and
7 whose corresponding group ID is equal to the first group ID value.

1 12. (Original): The method of claim 8 wherein the content-based code is a
2 hash code produced by applying a hash algorithm to content of a partition.

1 13. (Original): The method of claim 10 further comprising storing partition
2 identity information comprising a hash code and a group ID that correspond to each partition on
3 the storage system, wherein the first partition is associated with a first hash code value and a first
4 group ID value, wherein the step of producing a number of replicas includes adding information
5 which identifies each replica to the partition identity information, including the first hash code
6 value and the first group ID value.

1 14. (Original): The method of claim 8 wherein the first predetermined value
2 is less than the second predetermined value.

1 15. (Original): The method of claim 8 wherein the first predetermined value
2 is equal to the second predetermined value.

1 16. (Original): The method of claim 8 wherein the step of deleting one or
2 more replicas further includes deleting one or more replicas until all the replicas are deleted or
3 until the number of second partitions is less than a fourth predetermined value.

1 17. (Original): The method of claim 16 wherein the third predetermined value
2 is greater than the fourth predetermined value.

1 18. (Original): A method for accessing a storage system comprising:
2 receiving a request to store a file;
3 storing the file on the storage system;
4 identifying one or more partitions which collectively constitute the file, the
5 partitions referred to as input partitions;
6 storing partition information that is associated with the file, wherein the partition
7 information associates the file with each of its input partitions; and
8 for each input partition, if there are no identical partitions, then
9 if the number of replicas of the input partition is less than a threshold
10 value, then producing at least one replica of the input partition and storing the replica on
11 the storage system,
12 wherein an identical partition is a partition, other than the input partition,
13 of a file that is stored in the storage system whose content is identical to content of the
14 input partition.

1 19. (Original): The method of claim 18 wherein for a first input partition, if
2 there is at least one other file that comprises a partition that is identical to the first input partition,
3 then deleting a replica of the first input partition if the replica exists.

1 20. (Original): The method of claim 18 wherein for an input partition, the
2 number of identical partitions plus the number of replicas is equal to a predetermined value, the
3 threshold value being the difference between the predetermined value and the number of
4 identical partitions.

1 21. (Original): The method of claim 18 wherein the partition information
2 associated with the file comprises a partition ID for each input partition, wherein the partition ID
3 comprises a hash code and a group ID, wherein the hash code is determined by applying a hash
4 function to contents of a partition, wherein the group ID identifies a partition whose content is
5 unique among partitions which have the same hash code.

1 22. (Original): The method of claim 18 further comprising storing
2 information that identifies one or partition groups, a partition group comprising one or more
3 partitions identified from among one or more files which contain identical content, a partition
4 group further comprising one or more replicas of a partition in the partition group.

1 23. (Original): The method of claim 18 wherein the threshold value is one.

1 24. (Original): The method of claim 18 wherein the threshold value is a
2 number greater than one.

1 25. (Original): The method of claim 18 wherein each partition is the same
2 size as other partitions.

1 26. (Original): The method of claim 18 wherein identifying one or more
2 partitions includes determining a partition size by which the partitions of the file are identified.

1 27. (Original): The method of claim 18 wherein a partition size of partitions
2 of a file can be different for different files.

1 28. (Original): A data storage system comprising:
2 a storage component; and
3 a data processing component in data communication with the storage component,
4 the data processing component for receiving access requests from users, the access requests for
5 accessing data that is stored in the storage component or for storing data to the storage
6 component,
7 the data processing component configured to perform the method steps of:

8 accessing a first partition of a file, the first partition comprising a first
9 portion of data that constitutes the file;

10 if the first partition does not have a corresponding identical partition in the
11 storage component, then creating at least one replica; and

12 repeating for a second partition of the file, the second partition comprising
13 a second portion of the data.

1 29. (Original): The system of claim 28 wherein the data processing
2 component is further configured to perform the method steps of:

3 receiving a request to store data to the storage system, the request including the
4 data that constitutes the file; and

5 accessing the storage component to store the data.

1 30. (Original): The system of claim 28 wherein a partition is identified with a
2 partition ID, the partition ID being based on content of the partition, wherein partitions which
3 contain identical content have the same partition ID,

4 the storage system further comprising information which is stored in the storage
5 component, each partition having its associated partition identity information, the partition
6 identity information comprising a partition ID its associated partition, wherein partitions that are
7 identical have the same partition ID,

8 whereby identical partitions can be identified by consulting the partition identity
9 information.

1 31. (Original): The system of claim 30 wherein the partition ID comprises a
2 content-based code and a group ID, wherein the content-based code is determined from the
3 content of a partition, wherein if one partition and another partition have the same content-based
4 code but have different content, then the one partition is associated with a first group ID and the
5 other partition is associated with a second group ID different from the first group ID, wherein if
6 the one partition and the other partition have the same content-based code and have identical
7 content, then the one partition and the other partition both are associated with the same group ID.

1 32. (Original): The system of claim 31 wherein the content-based code is a
2 hash code, wherein the hash code is generated by applying a hash function to the content of a
3 partition.

1 33. (Currently amended): A data ~~processing~~storage system comprising:
2 first means for producing a partition ID for each partition comprising a file the
3 storage system, a partition comprising data from the file, the first means producing a first
4 partition ID for a first partition of a first file;
5 second means for identifying one or more identical partitions among other files in
6 the storage system based on a first partition ID; and
7 third means for creating a replica of the first partition in response to the second
8 means making a determination that there are no identical partitions,
9 wherein the first means, the second means, and the third means operate on every
10 partition comprising the first file.

1 34. (Original): The system of claim 33 wherein a partition ID comprises a
2 hash code and a group ID, the first means comprising hash means for producing a hash code
3 based on data comprising a partition and group ID means for determining a group ID, wherein

1 35. (Original): A method for accessing a storage system comprising:
2 accessing a first read-out partition of a data object, the first read-out partition
3 comprising at least a portion of content comprising the data object;
4 if content in the first read-out partition is corrupted, then:
5 accessing the storage system to find a replacement partition from among
6 one or more candidate partitions, including determining if a candidate partition is
7 corrupted or not, the replacement partition being a candidate partition that is not
8 corrupted; and
9 replacing content in the data object that constitutes the first read-out
10 partition with content of the replacement partition; and
11 repeating the foregoing with a second read-out partition of the data object.

1 36. (Original): The method of claim 35 wherein each partition is associated
2 with a partition ID, wherein if two partitions comprise identical content, then the two partitions
3 have the same partition ID, wherein the read-out partition is associated with a first partition ID
4 and the one or more candidate partitions each is associated with the first partition ID.

1 37. (Original): The method of claim 36 wherein the partition ID comprises a
2 hash code and a group ID, wherein the hash code is based on content of a partition, wherein if
3 the two partitions have the same hash code but their respective content is different from each
4 other, then the two partitions each have a different group ID, wherein if the two partitions
5 comprise identical content, then the two partitions have the same group ID.

1 38. (Original): The method of claim 35 wherein if a replacement partition
2 cannot be found for the first read-out partition, then indicating an error condition.

1 39. (Original): The method of claim 35 further comprising receiving a data
2 access request, the data access request including information indicative of the data object.

1 40. (Original): The method of claim 35 wherein the data object is a file.

1 41. (Original): A method for accessing data in a storage system comprising:
2 identifying a first data object;
3 obtaining a first partition of the first data object from the storage system;
4 performing a computation using data comprising the first data object to produce a
5 first computed value;
6 obtaining partition identification information relating to the first partition, the
7 partition identification information including a first previously computed value; and
8 if the first computed value does not match the first previously computed value,
9 then:

10 obtaining a first candidate partition from the storage system;
11 performing a computation using data comprising the first candidate
12 partition to produce a second computed value;

obtaining partition identification information relating to the first candidate partition, the partition identification information including a second previously computed value;

if the second computed value does not match the second previously computed value, then repeating with a second candidate partition; and

if the second computed value does match the second previously computed value, then replacing the data comprising the first partition of the first data object with the data comprising the first candidate partition.

42. (Original): The method of claim 41 wherein identifying a first data object is a step of receiving a read request for the first data object.

43. (Original): The method of claim 41 wherein if the first computed value does match the first previously computed value, then repeating with a second partition of the first data object.

44. (Original): The method of claim 41 wherein the first data object is a file.

45. (Original): The method of claim 41 wherein steps comprising the method are repeated for a second data object.

46. (Original): A data storage system comprising:
a storage subsystem; and
a data processing subsystem in data communication with the storage subsystem to store data to the storage subsystem and to access data stored on the storage subsystem, the data processing subsystem configured to:
access a first file stored on the storage subsystem, wherein the first file comprises data, the data being logically grouped into one or more accessed partitions;
determine, for each accessed partition, whether the accessed partition is corrupt, referred to as a corrupt partition;
determine, for each corrupt partition, whether there is a replacement partition on the storage system, the replacement partition being identical to the accessed partition at a time when the accessed partition was not corrupt; and

13. modify the first file to replace each of its corrupt partitions with a
14. replacement partition if it exists.

1 47. (Original): The system of claim 46 wherein the data processing subsystem
2 is further configured to communicate with one or more users and to receive access requests for
3 data stored on the storage subsystem, wherein one such access request is a request for the first
4 file.

1 48. (Original): The system of claim 46 further comprising partition identity
2 information, wherein each partition on the storage subsystem is associated with its corresponding
3 partition identity information, the partition identity information including a partition ID, wherein
4 the partition ID uniquely identifies content of a partition, wherein partitions that are identical
5 have the same partition ID,

6 wherein the data processing subsystem is further configured to determine whether
7 an accessed partition is corrupt based on its accessed content and on its partition ID.

1 49. (Original): The system of claim 46 wherein each partition on the storage
2 subsystem is associated with a partition ID comprising a hash code component and a group ID
3 component, wherein partitions that have the same hash code value also have the same group ID
4 if they have identical content, wherein the data processing subsystem is further configured to:
5 compute a first hash value for an accessed partition; and
6 compare the first hash value with the hash code component of the partition ID
7 associate with the accessed partition in order to determine whether the accessed partition is
8 corrupt.

1 50. (Currently amended): A data processing system comprising:
2 first means for accessing a partition comprising a target file that is stored on a
3 storage subsystem, a partition comprising data from the file, a partition having associated
4 therewith a partition ID that uniquely identifies content of the partition, wherein partitions
5 comprising identical content have the same partition ID;
6 second means for determining whether a partition is corrupt;

7 third means for identifying a replacement partition from among a plurality of
8 partitions stored on the storage subsystem to replace a corrupt partition, based on a partition ID
9 associated with the corrupt partition and on partition IDs of the plurality of partitions, the corrupt
10 partition being a constituent partition of a target file; and
11 fourth means to modify the target file to replace content comprising the corrupt
12 partition with content from a replacement partition.